

CARNEGIE-MELLON UNIVERSITY

DEPARTMENT OF COMPUTER SCIENCE

SPICE PROJECT

Spice Lisp User's Guide

Edited by Scott E. Fahlman and Monica J. Cellio

28 August 1984

Companion to the Mary Poppins Edition
of the Common Lisp Manual

Copyright © 1984 Carnegie-Mellon University

This is an internal working document of the Computer Science Department, Carnegie-Mellon University, Schenley Park, Pittsburgh, Pennsylvania 15213 USA . Some of the ideas expressed in this document may be only partially developed, or may be erroneous. Distribution of this document outside the immediate working community is discouraged; publication of this document is forbidden.

Supported by the Defense Advanced Research Projects Agency, Department of Defense, ARPA Order 3597, monitored by the Air Force Avionics Laboratory under contract F33615-81-K-1539. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Projects Agency or the U.S. Government.

Table of Contents

1. Introduction	1
1.1. Obtaining and Running Spice Lisp	1
2. Implementation Dependent Design Choices	3
2.1. Numbers	3
2.2. Characters	3
2.3. Vector Initialization	3
2.4. Packages	4
2.5. The Editor	4
2.6. Time Functions	4
2.7. Garbage Collection	4
3. Debugging Tools	5
3.1. Function Tracing	5
3.1.1. Encapsulation Functions	6
3.2. The Single Stepper	7
3.3. The Debugger	8
3.3.1. Movement Commands	9
3.3.2. Inspection Commands	9
3.3.3. Other Commands	10
3.4. Break Loop	11
3.4.1. Cleaning Up	12
3.5. Random Features	12
4. The Compiler	13
4.1. Calling the Compiler	13
4.2. Open and Closed Coding	13
4.3. Compiler Switches	14
4.4. Declare switches	15
5. Efficiency	16
5.1. Compile Your Code	16
5.2. Avoid Unnecessary Consing	16
5.3. Do, Don't Map	17
5.4. Think Before You Use a List	17
5.4.1. Use Vectors	17
5.4.2. Use Structures	18
5.4.3. Use Hashtables	18
5.4.4. Use Bit-Vectors	18
5.5. Simple Vs Complex Arrays	18
5.6. To Call or Not To Call	19
5.7. Keywords and the Rest	20

5.8. Numbers	20
5.9. Timing	21
6. Choosing Items From a Menu	22
6.1. The Menu Choose Functions	22
6.2. The Item	24
6.3. User Settable Attributes of the Choice Window	25
6.4. The Arrangement of Items in the Window	25
6.5. Some Examples	27
7. The Alien Data Type	29
7.1. The Alien-Structure Data Type	29
7.2. Defining Other Field Types	32
7.3. Variable-Format Structures	33
Index	34
Index	35

Chapter 1

Introduction

Common Lisp is a new dialect of Lisp, closely related to Maclisp and Lisp Machine Lisp. Common Lisp was developed in response to the need for a modern, stable, well documented dialect of Lisp that can be implemented efficiently on a variety of machine architectures.

Spice Lisp is the implementation of Common Lisp for microcodable personal machines running CMU's Spice computing environment. At present, Spice runs only on the Three Rivers Computer Corporation's PERQ; implementations for other machines are planned but not yet under way. Compatible versions of Common Lisp will soon be available for the DEC Vax, under both VMS and Unix, the Decsystem-20 with extended addressing, and the Symbolics 3600.

The central document for users of any Common Lisp implementation is the *Common Lisp Reference Manual*, by Guy L. Steele Jr. All implementations of Common Lisp must conform to this standard. However, a number of design choices are left up to the implementor, and implementations are free to add to the basic Common Lisp facilities. This document covers those choices and features that are specific to the Spice Lisp implementation. The *Common Lisp Reference Manual* and *Spice Lisp User's Guide*, taken together, should provide everything that the user of Spice Lisp needs to know.

For now, a number of documents describing useful library modules that run in Spice Lisp are included here. Once there are enough of these, the documents will be moved into a separate document on the Spice Lisp Program Library.

Spice Lisp is currently undergoing intensive tuning and development. For the next year or so, at least, new releases will be appearing frequently. This document will be modified for each major release, so that it is always up to date. Users of Spice Lisp at CMU should watch the SPICE and CLISP bulleting boards for release announcements, pointers to updated documentation files, and other information of interest to the user community.

1.1. Obtaining and Running Spice Lisp

In order to run Spice Lisp, you must have a Perq1a, Perq2, T1, or T2 with 16K control store. You must also have an up-to-date Accent system. Use the update program to get the current release of Accent. Then, decide

where you want the Spice Lisp files to live. There must be at least 5000 pages free in the partition you wish to put Spice Lisp in. It is suggested that you make a subdirectory called `slisp` in the user partition. Path to the directory you want to put Spice Lisp in, then run the update program with logical name `slisp-a`. When Spice Lisp is on your Perq, put the directory that it resides in on your search list and then just type `lisp` to the Accent shell.

Chapter 2

Implementation Dependent Design Choices

Several design choices in Common Lisp are left to the individual implementation. This chapter contains a partial list of these topics and the choices that are implemented in Spice Lisp.

2.1. Numbers

Currently, short-floats and single-floats are the same, and long-floats and double-floats are the same. Short floats use an immediate (non-consing) representation with 8 bits of exponent and a 21-bit mantissa. Long floats are 64-bit consed objects, with 12 bits of exponent and 53 bits of mantissa. All of these figures include the sign bit and, for the mantissa, the "hidden bit". The long-float representation conforms to the 64-bit IEEE standard, except that we do not support all the exceptions, negative 0, infinities, and the like.

Fixnums are stored as 28-bit two's complement integers, including the sign bit. The most positive fixnum is $2^{27} - 1$, and the most negative fixnum is -2^{27} . An integer outside of this range is a bignum.

2.2. Characters

Spice Lisp characters have 8 bits of code, 8 bits of font, and 8 control bits. The font bits are not used, and only 4 of the control bits are used: (control, meta, super, and hyper).

The control bit functions `Control`, `Meta`, `Super`, and `Hyper` are defined as in the *Common Lisp Manual*. The Perq keyboard does not produce these and Accent does not pass them to Spice Lisp, but programs can use these internally.

2.3. Vector Initialization

If no `:initial-value` is specified, vectors of Lisp objects are initialized to `nil`, and vectors of integers are initialized to 0.

2.4. Packages

Common Lisp requires four built-in packages: `lisp`, `user`, `keyword`, and `system`. In addition to these, Spice Lisp has separate packages `hemlock` and `hemlock-internals` (for the editor) and `compiler`.

2.5. The Editor

The `ed` function will invoke the Hemlock Editor. Hemlock is described in *The Hemlock User's Manual* and *The Hemlock Command Implementers Manual*; like Spice Lisp, it contains easily accessible internal documentation.

2.6. Time Functions

There is at present only one time function in Spice Lisp, due to the difficulty of getting such information from Accent.

`time` *form* [Macro]

The `time` macro evaluates its single *form* argument, prints the total *elapsed* time for the evaluation to `*trace-output*`, and returns the value which *form* returns.

2.7. Garbage Collection

Spice Lisp automatically does a GC whenever a user-specifiable ratio of stuff in dynamic space to available virtual memory is exceeded. This ratio is currently 2, meaning that when one has twice as much stuff in dynamic space as available virtual memory, a GC will be done. This initial setting is somewhat risky, since it assumes that half of the stuff in dynamic space is garbage. If you plan to make a lot of stuff that you want to keep around, you can set `lisp::gc-flip-ratio` to a smaller value, like 1 or 3/2. If you want to turn off GC entirely, you can set `lisp::*already-maybe-gcing*` to `t`.

Chapter 3

Debugging Tools

By Jim Large and Steve Handerson

3.1. Function Tracing

The tracer causes selected functions to print their arguments and their results whenever they are called. Options allow conditional printing of the trace information and conditional breakpoints on function entry.

`trace &rest specs` [Macro]

Invokes tracing on the specified functions,¹ and pushes their names onto the global list in `*traced-function-list*`. Each *spec* is either the name of a function, or the form

```
(function-name
  trace-option-name value
  trace-option-name value
  ...)
```

If no *specs* are given, then `trace` will return the list of all currently traced functions, `*traced-function-list*`.

If a function is traced with no options, then each time it is called, a single line containing the name of the function, the arguments to the call, and the depth of the call will be printed on the stream `*trace-output*`. After it returns, another line will be printed which contains the depth of the call and all of the return values. The lines are indented to highlight the depth of the calls.

Trace options can cause the normal printout to be suppressed, or cause extra information to be printed. Each traced function carries its own set of options which is independent of the options given for any other function. Every time a function is specified in a call to trace, all of the old options are discarded. The available options are:

- :condition** A form to eval before each call to the function. Trace printout will be suppressed whenever the form returns `nil`.
- :break** A form to eval before each call to the function. If the form returns non `nil`, then a breakpoint loop will be entered immediately before the function call.

¹Trace does not work on macros or special forms yet.

- `:break-after` Like `:break`, but the form is eval'd and the break loop invoked after the function call.
- `:break-all` A form which should be used as both the `:break` and the `:break-after` args.
- `:wherein` A function name or a list of function names. Trace printout for the traced function will only occur when it is called from within a call to one of the `:wherein` functions.
- `:print` A list of forms which will be evaluated and printed whenever the function is called. The values are printed one per line, and indented to match the other trace output. This printout will be suppressed whenever the normal trace printout is suppressed.
- `:print-after` Like `:print` except that the values of the forms are printed whenever the function exits.
- `:print-all` This is used as the combination of `:print` and `:print-after`.

`untrace &rest function-names` [Macro]
 Turns off tracing for the specified functions, and removes their names from `*traced-function-list*`. If no *function-names* are given, then all functions named in `*traced-function-list*` will be untraced.

`*traced-function-list*` [Variable]
 A list of function names which is maintained and used by `trace`, `untrace`, and `untrace-all`. This list should contain the names of all functions which are currently being traced.

`**trace-print-level*` [Variable]
`**trace-print-length*` [Variable]
`*print-level*` and `*print-length*` are bound to `*trace-print-level*` and `*trace-print-length*` when printing trace output. The forms printed by the `:print` options are also affected. `*Trace-print-level*` and `*trace-print-length*` are initially set to `nil`.

`*max-trace-indentation*` [Variable]
 The maximum number of spaces which should be used to indent trace printout. This variable is initially set to some reasonable value.

3.1.1. Encapsulation Functions

The encapsulation functions provide a clean mechanism for intercepting the arguments and results of a function.² `encapsulate` changes the function definition of a symbol, and saves it so that it can be restored later. The new definition normally calls the original definition.

²Encapsulation does not work for macros or special forms yet.

The original definition of the symbol can be restored at any time by the `unencapsulate` function. `Encapsulate` and `unencapsulate` allow a symbol to be multiply encapsulated in such a way that different encapsulations can be completely transparent to each other.

Each encapsulation has a type which may be an arbitrary lisp object. If a symbol has several encapsulations of different types, then any one of them can be removed without affecting more recent ones. A symbol may have more than one encapsulation of the same type, but only the most recent one can be undone.

`encapsulate` *symbol type body* [Function]

Saves the current definition of *symbol*, and replaces it with a function which returns the result of evaluating the form, *body*. *Type* is an arbitrary lisp object which is the type of encapsulation.

When the new function is called, the following variables will be bound for the evaluation of *body*:

`argument-list`

A list of the arguments to the function.

`basic-definition`

The unencapsulated definition of the function.

The unencapsulated definition may be called with the original arguments by including the form
(`apply basic-definition argument-list`)

`Encapsulate` always returns *symbol*.

`unencapsulate` *symbol type* [Function]

Undoes *symbol*'s most recent encapsulation of type *type*. *Type* is compared with `eq`. Encapsulations of other types are left in place.

`encapsulated-p` *symbol type* [Function]

Returns `t` if *symbol* has an encapsulation of type *type*. Returns `nil` otherwise. *Type* is compared with `eq`.

3.2. The Single Stepper

`step` *form* [Function]

Evaluates *form* with single stepping enabled or if *form* is `T`, enables stepping on until explicitly disabled. Stepping can be disabled by quitting to the lisp top level, or by evaluating the form (`step ()`).

While stepping is enabled, every call to `eval` will prompt the user for a single character command. The prompt is the form which is about to be eval'd. It is printed with `*print-level*` and `*print-length*` bound to `*step-print-level*` and `*step-print-length*`. All interaction is done through the stream `*query-io*`.

The commands are:

n (next)	Evaluate the expression with stepping still enabled.
s (skip)	Evaluate the expression with stepping disabled.
q (quit)	Evaluate the expression, but disable all further stepping inside the current call to step.
p (print)	Print current form. (does not use <code>*step-print-level*</code> or <code>*step-print-length*</code> .)
b (break)	Enter break loop, and then prompt for the command again when the break loop returns.
e (eval)	Prompt for and evaluate an arbitrary expression. The expression is evaluated with stepping disabled.
? (help)	Prints a brief list of the commands.
r (return)	Prompt for an arbitrary value to return as result of the current call to eval.
g	Throw to top level.

`*step-print-level*` [Variable]

`*step-print-length*` [Variable]

`*print-level*` and `*print-length*` are bound to these values when the current form is printed. `*Step-print-level*` and `*step-print-length*` are initially bound to some small value.

`*max-step-indentation*` [Variable]

Step indents the prompts to highlight the nesting of the evaluation. This variable contains the maximum number of spaces to use for indenting. Initially set to some reasonable number.

3.3. The Debugger

The debugger is an interactive command loop which allows a user to examine the active call frames on the Lisp function call stack. If it is invoked from an error breakpoint, it can show the function calls which led up to the error.

Inside the debugger, most commands refer to the *current stack frame*.³ The debugger assigns numbers to the frames on the stack, starting with zero as the most recent and increasing deeper into the stack. The debug prompt includes the number of the current frame as its main feature.

Most expressions typed to debug are simply evaluated as they would have been had you not entered debug. This includes the special debugger functions to be described, which are meaningful only inside the debugger. The biggest exception is the debugger commands, which are either one or two letters. These may display information about the current frame or change the current frame, but they generally do not affect the

³The debugging functions may refer to other frames by number. This will be described shortly.

evaluation history (*, **, and friends).

3.3.1. Movement Commands

These commands move to a new stack frame, and print out the name of the function and the values of its arguments in the style of a lisp function call. Frames which are not active are marked with a "***", and the reconstructed call consists of what arguments are present on the stack. `*Debug-print-level*` and `*debug-print-length*` affect the style of the printing.

Visible frames are those which have not been hidden by the `debug-hide` function which is described below. The special variable `*debug-ignored-functions*` contains a list of function names which are hidden by default.

- u** Move up to the next higher visible frame. More recent function calls are considered to be higher on the stack.
- d** Move down to the next lower visible frame.
- t** Move to the highest visible frame.
- b** Move to the lowest visible frame.
- f** Move to a given frame, visible or not. Prompts for the number.

3.3.2. Inspection Commands

These commands print information about the current frame and the current function.

- ?** Describe's the current function.
- a** Lists the arguments to the current function. The values of the arguments are printed along with the argument names.
- l** Lists the local variables in the current function. The values of the locals are printed, but their names are no longer available.
- p** Redisplays the current function call as it would be displayed by moving to this frame.
- pp** Redisplays the current function call using `*print-level*` and `*print-length*` instead of `*debug-print-level*` and `*debug-print-length*`.

(debug-value *symbol* [*frame*])

Returns the value of *symbol*, considered as a special variable, in the binding context of either the current frame, or the numbered frame, if specified.

(debug-local *n* [*Frame*])

Returns the value of the *n*th local variable in the current or specified frame.

(debug-arg *n* [*frame*])

Returns the *n*th argument of the frame.

(debug-pc [*frame*]) Returns the next instruction to be executed in the specified (active) frame. Can be used with Disassemble.

3.3.3. Other Commands

h Prints a brief but comprehensive list of commands on the terminal.

q Causes debug to return `nil`.

(debug-return *expression* [*frame*])

Forces the current function to return zero or more values. If the function was not called for multiple values, only the first value will be returned.

(backtrace) Prints a history of function calls. The printing is controlled by `*debug-print-level*` and `*debug-print-length*`. Only those frames which are considered visible by the frame movement commands will be shown.

(debug-hide *option* [*arg(s)*])

Makes the described stack frames invisible to the frame movement commands. The second argument is evaluated and may be a symbol or a list; the function returns the hidden members of the category. With no arguments, returns the current filter (hidden frames). *option* is a subcommand which may be one of:

package (*s*) Calls to hidden packages are visible, but calls within them are not.

function (*s*) Calls to the named functions will not be visible.

type (*s*) Hides miscellaneous frame and function types, any of:

compiled Calls to compiled functions will not be visible.

interpreted Calls to interpreted functions will not be visible.

lambdas Calls to lambda expressions will not be visible.

open Open frames will not be visible.

active Active frames will not be visible.

catch Catch frames will not be visible.

(debug-show *options arg or args*)

Cancels the effect of the corresponding `debug-hide`. Note that a frame may be hidden in a variety of other ways, though.

debug [Function]

Invokes the debugger. `debug` always returns `nil`.

debug-print-level [Variable]

debug-print-length [Variable]

print-level and ***print-length*** are bound to these values during the execution of some debug commands. When evaluating arbitrary expressions in the debugger, the normal ***print-level*** and ***print-length*** are in effect. These variables are initially set to some small number.

debug-ignored-functions [Variable]

A list of functions which are hidden by default. These functions can be made visible with the `debug` command `show`.

3.4. Break Loop

The break loop is a read-eval-print loop which is similar to the normal lisp top level. It can be called from any lisp function to allow the user to interact with the lisp system. When the user gives the command to exit the break loop, he may choose an arbitrary value for it to return.

When a lisp expression is typed in at the break loop's prompt, it is usually evaluated and printed. However, there are three special expressions which are recognized as break loop commands, and which are not evaluated.

\$G Typing this symbol causes a throw to the lisp top level: The current computation is aborted, and all bindings are unwound.

\$P Typing this symbol causes the break loop to return `nil`.

(`RETURN form`) Typing this expression causes the break loop to evaluate *form* and return the result(s).

The dollar sign character in the symbols **\$P** and **\$G** is intended to be the (escape) character -- ascii 27. For compatibility with the VAX VMS operating system, real dollar signs will be recognized also.

When the break loop is called, it tries to make sure that terminal interaction will be possible. All of the standard input output streams, ***standard-input***, ***standard-output***, ***error-output***, ***query-io***, and ***trace-output*** are bound to ***terminal-io*** for the duration of the break loop; and the state of the single stepper is bound to "off".

break tag &optional condition [Macro]

The break macro returns a form which prints the message "Breakpoint *tag*" to ***terminal-io*** and then invokes the break loop. If *condition* is present, then the form evaluates it and tests the result. If the result is `nil`, then the form returns `nil`; otherwise, the form prints the tag and invokes the break loop. *Tag* is never evaluated.

3.4.1. Cleaning Up

The break loop is called by the system error handlers. Since errors can happen unexpectedly, the break loop provides a mechanism for cleaning up any unusual state that a program may have caused.

****error-cleanup-forms****

[*Variable*]

A list of lisp forms which will be evaluated for side effects when a break loop is invoked. Whenever a break loop is entered, ***error-cleanup-forms*** will be bound to `nil`, and then the forms which were its previous value will be evaluated for side effects. There is no way to have the side effects undone when the break loop returns, and if any of the cleanup forms causes an error, the result can not be guaranteed.

As an example, a program that puts the terminal in an unusual mode might want to do something like this.

```
(let ((*error-cleanup-forms*
      (cons '(progn <code to restore terminal>)
              *error-cleanup-forms*)))
  <code to mess up terminal>
  .
  .
  .)
```

3.5. Random Features

describe object

[*Function*]

The **describe** function prints useful information about *object* on ***standard-output***. For any object, **describe** will print out the type. Then it prints other information based on the type of object. The types which are presently handled are:

- | | |
|-----------------|--|
| function | describe prints a list of the function's name (if any) and its formal parameters. If the name has documentation, then the documentation string will be printed. |
| symbol | The symbol's value, properties, and documentation are all printed. If the symbol has a function definition, then the function is described. |

The Compiler

4.1. Calling the Compiler

Functions may be compiled using `compile`, `compile-file`, or `compile-from-stream`. `Compile` operates exactly as documented in the *Common Lisp Reference Manual*.

```
compile-file &optional input-pathname &key :output-file :error-file      [Function]
          :lap-file :errors-to-terminal :load
```

This function is an expanded version of that described in the *Common Lisp Reference Manual*. If `input-pathname` is not provided `compile-file` prompts for it. `Output-file` and `Error-file` default to `T`, producing a fasl file and a compilation log with extensions `.fasl` and `.err`. `Lap-file` defaults to `nil`, indicating that the lap code should not be stored in a file. Any of these options may be `t`, `nil`, or the string name of a file to write to. *Errors-to-terminal* defaults to `T`; if specified and `nil` the compilation log goes only to the `.err` file. If `load` is specified and `non-nil` the compiled file is loaded after the compilation.

compile-from-stream input-stream [Function]

This function takes a stream as an input and reads lisp code from that stream until end of file is reached. The code is compiled and loaded into the current environment. No output files are produced.

4.2. Open and Closed Coding

When a function call is "open coded," inline code whose effect is equivalent to the function call is substituted for that function call. When a function call is "closed coded", it is usually left as is, although it might be turned into a call to a different function with different arguments. As an example, if `nthcdr` were to be "open coded" then

```
(nthcdr 4 foobar)
```

might turn into

```
(cdr (cdr (cdr (cdr foobar))))
```

or even


```
(do ((i 0 (1+ i))
      (list foobar (cdr foobar)))
    ((= i 4) list)).
```

If `nth` is "closed coded"

```
(nth x 1)
```

might stay the same, or turn into something like:

```
(car (nthcdr x 1)).
```

4.3. Compiler Switches

Several compiler switches are available which are not documented in the *Common Lisp Manual*. Each is a global special. These are described below.

peep-enable If this switch is non-nil, the compiler runs the peephole optimizer. The optimizer makes the compiled code faster, but the compilation itself is slower. ***peep-enable*** defaults to `t`.

peep-statistics
If this switch is non-nil, the effectiveness of the peephole optimizer (number of bytes before and after optimization) will be reported as each function is compiled. ***peep-statistics*** defaults to `t`.

inline-enable
If this switch is non-nil, then functions which are declared to be inline are expanded inline. It is sometimes useful to turn this switch off when debugging. ***inline-enable*** defaults to `t`.

open-code-sequence-functions
If this switch is non-nil, the compiler tries to translate calls to sequence functions into do loops, which are more efficient. It defaults to `t`.

optimize-let-bindings
If this is `t`, optimize some let bindings, such as those generated by lambda expansions and `self` based operations. If it is `:all`, optimize all lets. If it is `nil`, don't optimize any. It takes significant time to do all. The optimization involves replacing instances of variables that are bound to other variables with the other variables. Defaults to `t`.

examine-environment-function-information
If this is non-NIL, look in the compiler environment for function argument counts and types (macro, function, or special form) if you don't get the information from declarations. Defaults to `t`.

nthcdr-open-code-limit
This is the maximum size an `nthcdr` can be to be open coded. In other words, if `nthcdr` is called with `n` equal to some constant less than or equal to the ***nthcdr-open-code-limit***, it will be open coded as a series of nested `cdr`'s. ***nthcdr-open-code-limit*** defaults to 10.

complain-about-inefficiency

If this switch is non-nil, the compiler will print a message when certain things must be done in an inefficient manner because of lack of declarations or other problems of which the user might be unaware. This defaults to `nil`.

eliminate-tail-recursion

If this switch is non-nil, the compiler attempts to turn tail recursive calls (from a function to itself) into iteration. This defaults to `t`.

all-rest-args-are-lists

If non-nil, this has the effect of declaring every `&rest` arg to be of type list. (They all start that way, but the user could alter them.) It defaults to `nil`.

verbose

If this switch is `nil`, only true error messages and warnings go to the error stream. If non-nil, the compiler prints a message as each function is compiled. It defaults to `t`.

check-keywords-at-runtime

If non-nil, compiled code with `&key` arguments will check at runtime for unknown keywords. This is usually left on and defaults to `t`.

4.4. Declare switches

Not all switches for `declare` are processed by the compiler. The `f type` and `function` declarations are currently ignored.

The `optimize` declaration controls some of the above switches:

- ***peep-enable*** is on unless `cspeed` is greater than `speed` and `space`.
- ***inline-enable*** is on unless `space` is greater than `speed`.
- ***open-code-sequence-functions*** is on unless `space` is greater than `speed`.
- ***eliminate-tail-recursion*** is on if `speed` is greater than `space`.

Chapter 5

Efficiency

By Rob Maclachlan

In Spice Lisp on the Perq, as in any language on any computer, the way to get efficient code is to use good algorithms and sensible programming techniques, but to get the last bit of speed it is helpful to know some things about the language and its implementation. This chapter is a summary of various hidden costs in the implementation and ways to get around them.

5.1. Compile Your Code

In Spice Lisp, compiled code typically runs at least 100 times faster than interpreted code. Another benefit of compiling is that it catches many typos and other minor programming errors. Many lisp programmers find that the best way to debug a program is to compile the program to catch simple errors, then debug the code *interpreted*, only actually using the compiled code once the program is debugged.

Another benefit of compilation is that compiled (*ysfl*) files load significantly faster, so it is worthwhile compiling files which are loaded many times even if the speed of the functions in the file is unimportant.

Do Not be concerned about the performance of your program until you see its speed compiled -- some techniques that make compiled code run faster make interpreted code run slower.

5.2. Avoid Unnecessary Consing

Consing is the Lispy name for allocation of storage, as done by the `cons` function, hence its name. `cons` is by no means the only function which conses -- so does `make-array` and many other functions. Even worse, the Lisp system may decide to cons furiously when you do some apparently innocent thing.

Consing hurts performance in the following ways:

- Consing reduces your program's memory access locality, increasing paging activity.
- Consing takes time just like anything else.
- Any space allocated eventually needs to be reclaimed, either by garbage collection or killing your

Lisp.

Of course you have to cons sometimes, and the Lisp implementors have gone to considerable trouble to make consing and the subsequent garbage collection as efficient as possible. In some cases strategic consing can improve speed. It would certainly save time to allocate a vector to store intermediate results which are used hundreds of times.

5.3. Do, Don't Map

One of the programming styles encouraged by Lisp is a highly applicative one, involving the use of mapping functions and many lists to store intermediate results. To compute the sum of the square-roots of a list of numbers, one might say:

```
(apply #' + (mapcar #'sqrt list-of-numbers))
```

This programming style is clear and elegant, but unfortunately results in slow code. There are two reasons why:

- The creation of lists of intermediate results causes much consing (see 5.2).
- Each level of application requires another scan down the list. Thus, disregarding other effects, the above code would probably take twice as long as a straightforward iterative version.

An example of an iterative version of the same code:

```
(do ((num list-of-numbers (cdr num))
      (sum 0 (+ (sqrt (car num)) sum)))
    ((null num) sum))
```

Once you feel in you heart of hearts that iterative Lisp is beautiful then you can join the ranks of the Lisp efficiency fiends.

5.4. Think Before You Use a List

Although Lisp's creator seemed to think that it was for LIST Processing, the astute observer may have noticed that the chapter on list manipulation makes up less than ten percent of the COMMON LISP manual. The language has grown since Lisp 1.5, and now has other data structures which may be better suited to tasks where lists might have been used before.

5.4.1. Use Vectors

Use Vectors and use them often. Lists are often used to represent sequences, but for this purpose vectors have the following advantages:

- A vector takes up less space than a list holding the same number of elements. The advantage may vary from a factor of two for a general vector to a factor of sixty-four for a bit-vector. Less space means less consing (see 5.2).
- Vectors allow constant time random-access. You can get any element out of a vector as fast as you

can get the first out of a list if you make the right declarations.

The only advantage that lists have over vectors for representing sequences is that it is easy to change the length of a list, add to it and remove items from it. Likely signs of archaic, slow lisp code are `nth` and `nthcdr` -- if you are using these function you should probably be using a vector.

5.4.2. Use Structures

Another thing that lists have been used for is the representation of record structures. Often the structure of the list is never explicitly stated and accessing macros are not used, resulting in impenetrable code such as:

```
(rplaca (caddr (caddr x)) (caddr y))
```

The use of `defstruct` structures can result in much clearer code, one might write instead:

```
(setf (beverage-flavor (astronaut-beverage x)) (beverage-flavor y))
```

Great! But what does this have to do with efficiency? Since structures are based on vectors, the `defstruct` version would likewise take up less space and be faster to access. Don't be tempted to try and gain speed by trying to use vectors directly, since the compiler knows how to compile faster accesses to structures than you could easily do yourself. Note that the structure definition should be compiled before any uses of accessors so that the compiler will know about them.

5.4.3. Use Hashtables

In many applications where association lists (alists) have been used in the past, hashtables would work much better. An alist may be preferable in cases where the user wishes to rebind the alist and add new values to the front, shadowing older associations. In most other cases, if an alist contains more than a few elements, a hashtable will probably do the job faster. If the keys in the hashtable are objects that can be compared with `eq1` or better yet `eq`, then hashtable access will be speeded up by specifying the correct function as the `:test` argument to `make-hashtable`.

5.4.4. Use Bit-Vectors

Another thing that lists have been used for is set manipulation. In some applications where there is a known, reasonably small universe of items Bit-Vectors could be used instead. This is much less convenient than using lists, because instead of symbols, each element in the universe must be assigned a numeric index into the bit vector. Nevertheless, bit-vectors are *much* faster than lists.

5.5. Simple Vs Complex Arrays

Spice Lisp has two different representations for arrays, one which is accessed rapidly in microcode and one which is accessed much more slowly in Lisp code. The class of arrays which can be represented in the fast form corresponds exactly to the *one dimensional simple-arrays*, as defined in the COMMON LISP manual. Included in this group are the types `simple-string`, `simple-vector` and `simple-bit-vector`.

Declare Your Vector Variables -- If you don't the compiler will be forced to assume you are using the inefficient form of vector. Example:

```
(defun iota (n)
  (let ((res (make-vector n)))
    (declare (simple-vector n))
    (dotimes (i n)
      (setf (aref res i) i))
    res))
```

Warning: if you declare things to be simple when they are not, incorrect code will be generated and hard-to-find bugs will result. It is worthwhile to note, however, that system functions which create vectors will always create simple-arrays unless you force them to do otherwise.

5.6. To Call or Not To Call

The usual Lisp style involves small functions and many function calls; for this reason Lisp implementations strive to make function calling as inexpensive as possible. Spice Lisp is fairly successful in this respect. Function calling is not vastly more expensive than other instructions, and is certainly faster than procedure calling in Perq Pascal.

For this reason you should not be overly concerned about function-call overhead in your programs. *However*, function calling does take time, and thus is not the kind of thing you want going on in the inner loops of your program. Where removing function calling is desirable you can use the following techniques:

Write the code in-line

This is not a very good idea, since it results in obscure code, and spreads the code for a single logical function out everywhere, making changes difficult.

Use macros

A macro can be used to achieve the effect of a function call without the function-call overhead, but the extreme generality of the macro mechanism makes them tricky to use. If macros are used in this fashion without some care, obscure bugs can result.

Use inline functions

This is often the best way to remove function call overhead in COMMON LISP. A function may be written, and then declared inline if it is found that function call overhead is excessive. Writing functions is easier than writing macros, and it is easier to declare a function inline than to convert it to a macro. Note that the compiler must process first the inline declaration, then the definition, and finally any calls which are to be open coded for the inline expansion to take place.

Note that any of the above techniques can result in bloated code, since they have the effect of duplicating the same instructions many places. If code becomes very large, paging may increase, resulting in a significant slowdown. Inline expansion should only be used where it is needed. Note that the same function may be called normally in some places and expanded inline in other places.

5.7. Keywords and the Rest

COMMON LISP has very powerful argument passing mechanisms. Unfortunately, two of the most powerful mechanisms, rest arguments and keyword arguments, have a serious performance penalty in Spice Lisp.

The main problem with rest args is that the microcode must cons a list to hold the arguments. If a function is called many times or with many arguments, large amounts of consing may occur.

Keyword arguments are built on top of the rest arg mechanism, and so have all the above problems plus the problem that a significant amount of time is spent parsing the list of keywords and values on each function call.

Neither problem is serious unless thousands of calls are being made to the function in question, so the use of argument keywords and rest args is encouraged in user interface functions.

Another way to avoid keyword and rest-arg overhead is to use a macro instead of a function, since the rest-arg and keyword overhead happens at compile time and not necessarily at runtime. If the macro-expanded form contains no keyword or rest arguments, then it is perfectly acceptable to use keywords and rest-args in macros which appear in inner loops.

Note: the compiler open-codes most heavily-used system functions which have keyword or rest arguments, so that no run-time overhead is involved.

Optional arguments have no significant overhead.

5.8. Numbers

Spice Lisp provides five types of numbers for your enjoyment:

- fixnums bignums ratios short-floats long-floats

Only short-floats and fixnums have an immediate representation; the rest must be consed and garbage-collected later. In code where speed is important, you should use only fixnums and short-floats unless you have a real need for something else. Since `most-positive-fixnum` is more than one hundred million, you shouldn't need to use bignums unless you are counting the reasons to use Lisp instead of Pascal. Unfortunately the amount of floating point precision that will fit in twenty-eight bits is severely limited, so there are reasonable problems which require the use of long-floats.

Another feature of ratios and bignums which will keep you entertained for hours is that operations on these numbers are written in Lisp, not microcode; this results in orders of magnitude slower execution.

Printing of long-floats is painfully slow -- around three seconds. While you wait, consider that the float printing algorithm is the only known correct float printing method. Other methods run in real time, but they

lose precision in the low-order digits.

5.9. Timing

One good way of improving a program's performance is to make extensive timings to find code which is time-critical. Spice Lisp contains one time function, `time`, which finds total elapsed time. For things which execute fairly quickly it may be wise to time more than once, since there may be paging overhead in the first timing. The times that `time` gets are only accurate to a certain number of decimal places, so for small pieces of code it may be a good idea to write a *compiled* driver function which calls the function to be tested a few hundred times. If one finds the time and divides by the number of iterations, then fairly accurate statistics can be collected.

Chapter 6

Choosing Items From a Menu

By Jim Müller

Spice Lisp provides a set of routines for creating and using menus. This chapter describes these routines.

6.1. The Menu Choose Functions

`menu-prepare` *items position-x position-y in-window stay-in-place title title-font default-item*
ncolumns items-justified side-margin top&bottom-margin spacing font label-font abort-value [Macro]

`menu-choose-from-structure` *menu* &optional *paste-up* [Function]

`Menu-prepare` takes a list of items to display in a pop up window on the screen and various other information; it returns a structure with the given information included. It is intended that this structure later be given to `menu-choose-from-structure`, and at that time the user may mouse a desired item from the list, causing a value (specified with the item in the call to `menu-prepare`) to be returned. Labels and column break markers may be embedded in the list of items; the person calling the function may specify display information, such as fonts, for each individual item or for all items. The items list is not evaluated, but the font information and item names are evaluated (in the call's lexical environment, at call time) as are the return values (in the call's lexical environment, at mouse click time, when the call to `menu-choose-from-structure` is made). A help string may be supplied with each item; it is not evaluated.

Position-x and *position-y* default to the pixel x and y coordinates of the mouse pointer, and *in-window* defaults to the window encompassing the whole screen. These three variables specify where the menu should pop up. If *stay-in-place* is specified and not `nil` the window will appear with its upper left hand corner at exactly that position, or an error will be signalled if this is not possible. Otherwise it may be moved in order to make it fit on the screen.

The *title* is a string to be displayed above the rest of the menu in inverse video in *title-font*. If no title is given, the inverse video title bar does not appear.

Paste-up is accepted by `menu-choose` and `menu-choose-from-structure`. If it is t the

menu is put constructed and on the screen, but the user does not choose from it. If it is `:and-choose` everything proceeds as normal but the menu stays on the screen after the mouse-click. In either case the user may later get to the menu using `pasted-menu-choose`. Those menus which are pasted on the screen are stored in `*pasted-menus*`.

The *default-item* is the name of the item on which the mouse initially appears. The default default-item is the item in the middle of the list. If a bogus default item is specified, it is ignored. If a menu structure is created, the default item is the last one chosen from it.

Ncolumns specifies how many columns will be used. Its value defaults to one. See the section below on arrangement of items in the window.

Spacing specifies the number of pixels width between columns. This much white space is inserted after all columns but the last. *Spacing* defaults to twice the width of the letter x in *font*. *Items-justified* tells whether to center the items, or justify them to the left or right (`:center`, `:left`, `:right`); this defaults to `:center`.

Top&bottom-margin is the number of pixels of margin left at the top and bottom of the menu. *Side-margin* acts in the obvious corresponding way. Both default to ten.

Font is the default font for items which are not labels to appear in, and defaults to whatever the font normally defaults to. *Label-font* is the default font for labels to appear in, and defaults to the italic font corresponding to the default font.

Abort-value is a value which is returned if the user types control-g while choosing from a menu; it defaults to `nil`.

If the optional *paste-up* argument to `menu-choose-from-structure` is supplied, the menu is simply put in a permanent place on the screen; the user does not get to choose from it. If *paste-up* is `:and-choose` it is permanently pasted on the screen and the user gets to choose as well.

`menu-choose &key items position-x position-y in-window stay-in-place title title-font default-item
ncolumns items-justified side-margin top&bottom-margin spacing font label-font abort-value
paste-up` [Macro]

This combines `menu-prepare` and `menu-choose-from-structure` into one call. The *paste-up* argument is passed to the latter function, all others to the former. `menu-choose` returns the value which `menu-choose-from-structure` returns.

`pasted-menu-choose menu` [Function]

This takes either a menu structure which is already on the screen or the title of such a menu, and allows the user to choose from it. `pasted-menu-choose` returns `nil` if the argument is a title or a menu that is not pasted up or does not exist, and signals an error for any other argument.

`unpaste-menu menu` [Function]

This takes a menu structure or which is pasted onto the screen (or its title), removes it, and returns `nil`.

`pasted-menu-p menu` [Function]

This predicate tells whether its argument is a menu that is currently pasted on the screen, or the title of one.

`Pasted-menu-choose` takes may take a menu prepared by `menu-prepare`, a menu which has been pasted up by some means, or the title of a menu which has been pasted up by some means.

6.2. The Item

Each item that the user may choose consists of a string (or character or symbol) to display and some other information. The additional information may tell how the string should be printed, what to return if the user selects the item, or possibly a note saying that the item is just for show and may not be chosen.

The item may be just a name x , in which case x is displayed and (string x) is returned if the item is chosen. An item may be the keyword `:new-column`, which is not really displayed at all but signals that a column break is desired, for easy use of multiple columns. Finally, an item may be a list, whose `car` is a name, and whose `cdr` is a bunch of keyword arguments telling about the item. In fact each of these keyword arguments is just zero or more forms, giving the item the following format:

```
item ::= name | :new-column | (name {keyword arg*}*)
```

```
name ::= symbol | string | character.
```

In fact, each keyword knows how many args it wants. If it wants a particular number it simply grabs that many; if it wants a variable number then it may scan for the next keyword as a cue that it has seen all of its args. If at any time not enough args are left for a keyword, or it is time to start parsing a new keyword arg and a keyword does not follow, an error is signalled.

The following sorts of keyword arguments may be provided:

- `:help` string. The string is given as help in the lisp window title line as long as the entry *name* with help *help* is being moused.
- `:value` value. The value is returned if the user clicks the mouse button on the item.
- `:values` ({value}*). The values are returned as multiple values if the user clicks the mouse button on the item.
- `:eval` form. The form is evaluated and returned when the user clicks the mouse button on the item.

- `:funcall fn ({arg}*)`. The result of the obvious funcall is returned when the user clicks the mouse button on the item.
- `:no-select`. The item may not be selected. Clicking the mouse here just causes the lisp window to flash.
- `:buttons {button-specifier}*`. This takes some button specifiers, each of which is a list. The specifier's car specifies a particular button, `:left`, `:middle`, or `:right`. The rest of the specifier list is a keyword arg of the `:values`, `:eval`, `:funcall`, or `:no-select` form; if there is anything after the first keyword arg it is ignored. Thus each specifier says what to return if that button is pressed, or that that button is inactive while mousing this item. If a button is left unspecified it causes `nil` to be returned. A button may not be specified more than once.
- `:font font`. This specifies the font in which the item name is to be printed. Currently the only font available is the default one, but this should change soon.
- `:label`. This is an abbreviation for `:font label-font :no-select`.

Help, font, and inverse are really the only options which may be meaningfully combined with other options. If conflicting things are specified, the one most recently specified takes precedence, so that `("Hosts" :label :font bold)` causes "Hosts" to be displayed as a label, boldface and not selectable, and `("X" :eval (if t 3) :values 'cmu-cs-spice)` returns cmu-cs-spice if selected.

6.3. User Settable Attributes of the Choice Window

Some of the parameters to `menu-choose` and `menu-prepare` default to global variables whose values are initialized to the values stated in the first section. These parameters are controlled this way because they are relatively independent of the other parameters and are the sorts of things users may have preferences about. Of course, if the function caller specifies a value for the parameter, that value is used.

Position-x, *position-y*, and *in-window* are controlled by `*default-menu-choose-position-x*`, `*default-menu-choose-position-y*`, and `*default-menu-choose-in-window*`. *Ncolumns* is controlled by `*default-menu-choose-ncolumns*`. *Items-justified* is controlled by `*default-menu-choose-items-justified*`.

6.4. The Arrangement of Items in the Window

The shape of the window may be specified by putting `:new-column` markers in the items list. In this case the function simply does as told and its only intervention is when there are too many columns to fit in the allowed width, in which case it signals an error.

If no `:new-column` markers are present the window shape is controlled by specifying the parameter *ncolumns*. This parameter must be a positive integer number of columns (less than the number of items) to use, and it defaults to one.

If the menu is too wide to fit on the screen, or some columns are too tall, an error is always signalled.

6.5. Some Examples

Following is the initial display (resulting near the mouse) from the call:

```
(let ((cmuc (make-machine 'cmuc)))
  (menu-choose :items
    '(("Machines" :label :help "Important sites at CMU-CSD")
      ("Spice" :values ("Vax" 780) :help "The Spice Vax.")
      (C :buttons
         :funcall #'logout-everyone (cmuc)
         :value "CmuC"
         :funcall #'(lambda (x) (login-on-machine x cmuc)) (*me*)
         :help
          "Left button clears the C of non-slispsers. Middle button just returns the
          real machine name. Right button logs me onto the C.")
        ("Lisp2" :eval (full-name (get-current-user 'lisp2))
          :help "Returns the full name of the current hacker on the rack Perq2.")
        ("Beryl" :font (bold-font current-font) :values (t t t)
          :help "Alto, downtown/Forbes corner of the rack.")))
    :label-font (bold-italic-font current-font)))
```

Machines
 Spice
 C
 Lisp2
 Beryl

Note that the word machines is not a title.

The following initial display results in the upper left hand corner of the lisp window from the call:

```
(setq baz
(menu-prepare :items
'(("Slispers" :label :help "The good people")
  ("Scott Fahlman" :help "Santa Claus")
  ;; With a name like this, obviously beyond help.
  "Joseph Wholey"
  ("Rob MacLachlan" :font 'peanut12 :help "Ross Thompson imitator.")
  ("Jim Muller" :no-select)
  :new-column
  (|Dark Siders| :label :help "THE BAD PEOPLE")
  (rashid :values (t 6.4 (cons 3 #\})) :help "ACCENT")
  (nason :funcall #' + (v-value w-value))
  ("ED SMITH" :help "LAUNCH")))

:default-item "Joseph Wholey" :spacing (font-string-width "xxxxxxx" current-font)
:title "Frob" :font (italic-font current-font) :position-x 0 :position-y 0))

(menu-choose-from-structure baz)
```

	Frob	
<i>Slispers</i>		<i>Dark Siders</i>
<i>Scott Fahlman</i>		<i>RASHID</i>
<i>Joseph Wholey</i>		<i>NASON</i>
		<i>ED SMITH</i>
<i>Jim Muller</i>		

Note that Rob's name doesn't appear, since the Perq doesn't display most characters in peanut12.

Chapter 7

The Alien Data Type

By Jim Large and Dan Aronson

A problem that arises in many Common Lisp implementations is dealing with the complex structured records or messages that are exchanged at the interface between Lisp and the outside world. Such alien data structures will typically be collections of integers, floating point numbers, strings, boolean flags, bit vectors, enumerated types represented as small integers, and so on. All of these types have some rough correspondence with internal Lisp data-types, but at the time of their arrival and departure they will be in whatever implementation-dependent format is expected by the alien software, and the Lisp garbage collector must treat the alien object as an unstructured vector of bits or bytes.

Given a knowledge of the structure of an alien record, it is relatively easy for the Lisp-level code to convert each field of the message into the corresponding Lisp form, but we want this knowledge to be concentrated in one place so that changes in the external message format can be easily accommodated by the Lisp code. What is needed is a convenient form for specifying how the alien record is to be parsed and packed and how each of its fields is to be interpreted as a Lisp object. In a manner similar to `defstruct`, this specification will be processed to create a family of field-accessing and field-altering macros that perform the proper translations, in addition to doing the access. Thus, once the structure of the alien record has been specified, it is no harder to access than the fields of a `defstruct`.

This new facility is built into Vax Common Lisp and Spice Lisp.

7.1. The Alien-Structure Data Type

There is a new data-type called `alien-structure`. This is just a new structure-type defined by `defstruct`. The alien structure contains a name (a Lisp symbol), a length (number of 8-bit bytes in the data vector), and a pointer to the actual blob of uninterpreted bits. In Spice Lisp and Vax Common Lisp, this blob is an packed-fixnum vector (a U-Vector, in our internal parlance) of 8-bit bytes; other implementations might have to use a bit-vector for this. One could ask `typep` if an object is an `alien-structure`, and could access the innards via `alien-structure-name` and `alien-structure-data`. One can also find the length of the data area (in bytes) using the macro `alien-structure-length`.

Alien structures are created by a macro, `def-alien-structure`, that parallels `defstruct` in form:


```
(def-alien-structure (name option1 option2 ...)
  field-description-1
  field-description-2
  ...)
```

where the field descriptions are of the form

```
(field-name alien-field-type start end
  field-option-name-1 field-option-value-1
  field-option-name-2 field-option-value-2
  ...)
```

The options in `def-alien-structure` are a subset of those allowed in `defstruct`: `:conc-name`, `:constructor`, `:predicate`, `:print-function`, and `:eval-when`. Alien structures are always named and the user cannot specify a `:type` option or any of the array options.

In addition, there is a `:length` option that takes as its argument the length of the data area in 8-bit bytes. This is used when new instances of this structure-type are created within Lisp by the constructor macro. If `:length` is not specified, the length defaults to the maximum of the *end* values of the fields making up the structure, ignoring any fields with *end* values of `nil`. When alien structures are read in from outside the Lisp, `:length` controls the allocated length of the data vector, but the actual length (as reported by `alien-structure-length`) is set by the size of the incoming block of data. An error will be signalled if the incoming data block does not fit within the allocated length. If no `:length` is specified in this case, the default is to allocate a vector the same size as the incoming data block.

Each field has a name, perhaps modified by `:conc-name`. The field options are `:invisible` and `:read-only`, as in `defstruct`, plus

`:default-value`. The latter is a value that is placed into the slot at the time the alien structure is created, if no other value is specified at that time. This value is inserted into the alien structure as if by `setf`, and it is therefore processed by the field-type conversions (see below). If no `:default-value` is specified, the field is initialized to 0. (This initialization is only done if the data block is created within the Lisp; if it arrives from outside, the bits are left alone unless specifically altered by the user.)

The *start* and *end* values for a field indicate where, in the alien structure's data area, the field is to be found. These numbers are in terms of bytes. As usual, they are zero-based, *start* is inclusive and *end* is exclusive. If the *end* value is `nil`, the field has no fixed length, but runs from the specified *start* to the *end* of the data block, as indicated by `alien-structure-length`. When such a field is written into, the `alien-structure-length` will be adjusted to reflect the new *end* of the field; however, any attempt to extend the field beyond the allocated length of the data vector will signal an error.

It is possible for two fields to overlap; sometimes this will be useful when one field wants to be interpreted in two different ways. Obviously, if two overlapping fields are written into, the later write clobbers the results of the earlier one. It is also possible to have gaps between the defined fields; these would correspond to parts of an incoming message that are uninteresting to the Lisp program, for example. Such gaps are initialized to 0

when the alien structure is created within the Lisp, unless the block of data comes in from outside. If the block comes from the outside, the bits in the inter-field gaps are not altered. Fields may appear in any order within `def-alien-structure`.

In the rare cases where the boundaries of a field do not land on byte-boundaries, rational numbers may be supplied as the *start* and *end* values. So most of the time you can pick up a string from (16 32) or an integer from (0 2), but sometimes you would get a boolean from (3/8 4/8) or an integer from (1/2 3/2). An error is signalled if the rational does not specify an integral bit-address.

The alien-field-type argument is a symbol that tells how the field is to be interpreted by the Lisp system. Each alien field type associates a particular alien-format representation with some internal Lisp data-type; functions exist for turning the contents of the field into the internal object and for packing an internal object of the right type back into the field. Some of these types will be built-in:

- string** On access, the field is interpreted as a string, one character per byte, and the corresponding Lisp string is returned. The `setf` form accepts a Lisp string and puts the characters into the field.
- perq-string** Like `string` except that the first byte in the field is the number of characters, and the remaining bytes are the characters. On access, the length of the Lisp string will be determined by the first byte of the field. The `setf` form will set the first byte according to the length of the Lisp string. The size of the field may not be greater than 256 bytes.
- signed-integer** The field is interpreted as a signed, two's complement integer, and the corresponding Lisp integer is returned. On write, the process is reversed.
- unsigned-integer** The field is interpreted as an unsigned, positive integer, and the corresponding Lisp integer is returned.
- bit-vector** The field is returned as a bit-vector.
- port** The field contains an Accent IPC port specified as a 32-bit integer. The internal Lisp format is a `port` structure, with the integer value hidden inside.
- (selection *s0 s1 s2 ...*)** The *S_n* are evaluated (at access or `setf` time) to produce arbitrary Lisp objects. On access the alien field is interpreted as an unsigned integer, and the corresponding *S_n* value is returned inside the Lisp. On output, the setting function receives one of the values and stores the corresponding integer into the field. Comparison of items against *S_n* values is done with `eq1`.
- ieee-single-flonum** The field is interpreted as containing a 32-bit IEEE-format flonum, and this is returned as the internal Lisp flonum type that most closely matches this type. This will vary from one implementation to another, but will be constant within a given implementation. An

implementation would provide whatever floating formats are important in its host environment -- the VAX might provide D, F, G, and H formats rather than IEEE formats.

In each case, the `setf` form will signal an error if the specified field is too small to hold the item coming from Lisp. For integers, the error will occur if significant bits would be lost in doing the write.

7.2. Defining Other Field Types

In addition to these built-in primitive alien field types, the user can define his own via `def-alien-field-type`, a macro with the following arguments:

```
(def-alien-field-type name internal-type primitive-type
                      access-fn setf-fn)
```

internal-type is a Common Lisp type specifier indicating the type of internal Lisp object that the field will be mapped to. *primitive* is any pre-defined alien-field-type: one of the primitives defined above or a field type defined earlier by `def-alien-field-type`. On access, this is applied to the alien object to extract a Lisp object; then this object is passed to the access function, usually a function of one argument, for further processing. For a `setf`, the new value is first passed through the `setf` function, also usually a function of one argument; the result of this is then packed into the alien structure as indicated by *primitive-type*.

For example, suppose we wanted to create a new field-type named `backwards-string`, in which the alien field is treated as a reversed string. This would be done as follows:

```
(def-alien-field-type reverse-string
  'string ; Turns to string internally.
  'string ; Primitive access to get a string.
  #'(lambda (x) (reverse x)) ; Reverse string on access.
  #'(lambda (x) (reverse x)) ; Also reverse string on setf.
```

Once this is done, the *reverse-string* field type can be used in `def-alien-structure-type` and as a primitive in defining still more complex field types.

Sometimes, it is desirable to create an alien field type in which the access and `setf` conversion functions can take additional parameters. The *selection* field-type, listed above among the primitives, is one such type. To achieve this effect, one defines an alien field type whose access and `setf` functions take more than one argument. The additional arguments should be optional. When a field-type expression in `def-alien-field` is a list rather than a symbol, the `car` of the list is the type name, and the remaining elements are expressions which are evaluated at access and `setf` time. The results of these evaluations are passed to the access and `setf` functions as additional arguments. This all sounds more complex than it really is. To produce the *selection* field type, if this were not built in as a primitive, one would do the following:

```

(def-alien-field-type selection
  't ; Produces any kind of Lisp object.
  'unsigned-integer ; Primitive access as unsigned integer.
  #'(lambda (n &rest s-list)
      (nth n s-list)) ; Select Nth value in list of choices.
  #'(lambda (x &rest s-list)
      (position x s-list))) ; Find index of item in list, EQL test.

```

7.3. Variable-Format Structures

The machinery described above is optimized for dealing with alien structure types whose fields are fixed in size and position at the time the structure-type is defined. Given the nature of software outside the Lisp world, this is the sort of thing we will be seeing the most of. However, it would also be nice to be able to use the alien field type translation for packing and unpacking variable-format records. To handle this variable case without getting too complicated, the following simple packing and unpacking macros are provided:

```
(alien-field alien-structure alien-field-type start end)
```

This form can be used to access an arbitrary field in any type of alien structure, using the specified alien-field-type, *start*, and *end*. This form pays no attention to any fixed-position fields that may have been defined for structures of this type; it just does what you tell it to do. The alien-field-type may be any pre-defined type. The *start* and *end* arguments are expressed in terms of 8-bit bytes, and may be ratios if it is necessary to reference a field that does not lie on even byte boundaries. This form may be used within a *setf* to alter a field.

```

(pack-alien-structure name length
  (value alien-field-type start end)
  (value alien-field-type start end)
  ...)

```

This creates and returns a new alien-structure object with the specified name and length (in bytes). The name may or may not be an alien-structure-type that has already been defined; in any event, the name is simply stored away and has no effect on how this object is filled. Each of the values is evaluated (to produce a Lisp object) and then is packed into the specified place in the new alien structure using the *setf*-transform of the specified alien field type. For example, to send a variable-length string from Lisp to wherever, something like the following function might be employed:

```

(defun export-string (s)
  (send-external-message wherever
    (pack-alien-structure 'string-message (+ (length s) 4)
      ((length s) 'unsigned-integer 0 4)
      (s 'string 4 (+ 4 (length s))))))

```

To receive a string of the same format:

```

(defun import-string ()
  (let* ((foo (receive-external-message wherever))
    (string-length (alien-field foo 'integer 0 4)))
    (alien-field foo 'string 4 (+ string-length 4))))

```

Index

Index

****error-cleanup-forms**** variable 12
****trace-print-length*** variable 6
****trace-print-level*** variable 6

break macro 11

compile-file function 13
compile-from-stream function 13

debug function 11
debug-ignored-functions variable 11
debug-print-length variable 11
debug-print-level variable 11
describe function 12

encapsulate function 7
encapsulated-p function 7
:error-file keyword
 for compile-file 13
:errors-to-terminal keyword
 for compile-file 13

:lap-file keyword
 for compile-file 13
:load keyword
 for compile-file 13

max-step-indentation variable 8
max-trace-indentation variable 6
menu-choose macro 23
menu-choose-from-structure function 22
menu-prepare macro 22

:output-file keyword
 for compile-file 13

pasted-menu-choose function 23
pasted-menu-p function 24

step function 7
step-print-length variable 8
step-print-level variable 8

time macro 4
trace macro 5
traced-function-list variable 6

unencapsulate function 7
unpaste-menu function 24
untrace macro 6